

AD-A222 395

LEARNING EVENTS IN THE
ACQUISITION OF THREE SKILLS

Technical Report PCG - 20

Kurt VanLehn
Departments of Psychology & Computer Science

DEPARTMENT
of
PSYCHOLOGY



DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

Carnegie Mellon University

DTIC
ELECTE
JUN 06 1990
S E D

90 06 05 054

LEARNING EVENTS IN THE ACQUISITION OF THREE SKILLS

Technical Report PCG - 20

Kurt VanLehn
Departments of Psychology & Computer Science

Carnegie Mellon University
Pittsburgh, PA 15213 U.S.A.

July 21, 1989

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



In G. Olson & E. Smith (Eds.) *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
Also available as Technical report AIP - 73

This is the final report on the research supported by the Personnel and Training Research Program, Psychology Sciences Division, Office of Naval Research, under contract N00014-86-K-0349. Reproduction in whole or in part is permitted for any purpose of the United States Government. Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; Distribution unlimited	
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE				
4. PERFORMING ORGANIZATION REPORT NUMBER(S) PCG - 20			5. MONITORING ORGANIZATION REPORT NUMBER(S) Same as Performing Organization	
6a. NAME OF PERFORMING ORGANIZATION Carnegie Mellon University		6b. OFFICE SYMBOL (If applicable)		7a. NAME OF MONITORING ORGANIZATION Personnel and Training Research Office of Naval Research (Code 1142PT)
6c. ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, Pennsylvania 15213			7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, VA 22217-5000	
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Same as Monitoring Organization		8b. OFFICE SYMBOL (If applicable)		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-88-K-0086
8c. ADDRESS (City, State, and ZIP Code)			10. SOURCE OF FUNDING NUMBERS	
			PROGRAM ELEMENT NO N/A	PROJECT NO. N/A
			TASK NO. N/A	WORK UNIT ACCESSION NO N/A
11. TITLE (Include Security Classification) Learning Events in the Acquisition of Three Skills				
12. PERSONAL AUTHOR(S) VanLehn, Kurt Alan				
13a. TYPE OF REPORT Technical		13b. TIME COVERED FROM 88Jan01 to 90Dec31		14. DATE OF REPORT (Year, Month, Day) 1989 July 21
15. PAGE COUNT 8				
16. SUPPLEMENTARY NOTATION In G. Olson & E. Smith (Eds.) "Proceedings of the Eleventh Annual Conference of the cognitive Science Society," Hillsdale, NJ: Erlbaum. Also AIP-73.				
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	strategy discovery, skill acquisition, machine learning, cognitive science, impasses-driven learning,	
19. ABSTRACT (Continue on reverse if necessary and identify by block number)				
<p>➤ According to current theories of cognitive skill acquisition, new problem solving rules are constructed by proceduralization, production compounding, chunking, syntactic generalization, and a variety of other mechanisms. All these mechanisms are assumed to run rather quickly, so a rule's acquisition should be a matter of a few seconds at most. Such "learning events" might be visible in protocol data. This paper discusses a method for locating the initial use of a rule in protocol data. The method is applied to protocols of subjects learning three tasks: a river crossing puzzle, the Tower of Hanoi, and a topic in college physics. Rules were discovered at the rate of about one every half hour. Most rules required several learning events before they were used consistently, which is not consistent with the one-trial learning predicted by explanation-based learning methods. Some observed patterns of learning events were consistent with a learning mechanism based on syntactic generalization of rules. Although most rules seem to have been acquired at impasses - occasions when the subject does not know what to do next - there were clear cases of rules being learned without visible signs of an impasse, which does not support the popular hypothesis that all learning occurs at impasses.</p>				
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL Susan Chipman			22b. TELEPHONE (Include Area Code) (202) 696-4322	22c. OFFICE SYMBOL 1142 PT

Learning events in the acquisition of three skills

Kurt VanLehn

Departments of Psychology and Computer Science
Carnegie-Mellon University

Abstract

According to current theories of cognitive skill acquisition, new problem solving rules are constructed by proceduralization, production compounding, chunking, syntactic generalization, and a variety of other mechanisms. All these mechanisms are assumed to run rather quickly, so a rule's acquisition should be a matter of a few seconds at most. Such "learning events" might be visible in protocol data. This paper discusses a method for locating the initial use of a rule in protocol data. The method is applied to protocols of subjects learning three tasks: a river crossing puzzle, the Tower of Hanoi, and a topic in college physics. Rules were discovered at the rate of about one every half hour. Most rules required several learning events before they were used consistently, which is not consistent with the one-trial learning predicted by explanation-based learning methods. Some observed patterns of learning events were consistent with a learning mechanism based on syntactic generalization of rules. Although most rules seem to have been acquired at impasses—occasions when the subject does not know what to do next—there were clear cases of rules being learned without visible signs of an impasse, which does not support the popular hypothesis that all learning occurs at impasses.

Introduction

The goal of this research is to see if people's behavior during the initial construction of a rule sheds any light on which of the many contemporary models of rule acquisition is a better characterization of human learning methods. Many theories of cognitive skill acquisition assume that rules are initially formed during events lasting only a few seconds. This is about the same time scale as a protocol line, so such "learning events" might be visible in a protocol. However, it is difficult to locate the exact line of a protocol where a rule is formed. All the existing simulations of cognitive skill acquisition (e.g., Anderson, Farrell & Saurers, 1984; Anzai & Simon, 1979; Klahr, Langley & Neches, 1987; VanLehn, 1983, in press) use multi-line episodes or other units of analysis that are too large for this purpose.

This paper first presents a method for locating protocol lines where rules are acquired, then discusses the results obtained by applying it to protocols from three task domains. Two of the tasks are puzzles: the Tower of Hanoi and a river-crossing puzzle. Although these two studies were intended merely to check that the analytic method works, they nonetheless yielded some interesting findings. In the third study, the task domain is college physics. This simulation is still under construction, so only some preliminary results can be reported. All three tasks are similar in that the subjects are "learning by doing" — they are solving problems without help from a tutor or an instruction manual. It is an open question whether the results and methods discuss herein extend to other instructional situations.

A method for locating protocol lines where rules are learned

The following two-step method is used to locate lines of the protocol corresponding to the construction of a rule. The first step is to postulate a large set of plausible rules for problem solving in the task domain. The rules can be inferred from analyzing other subjects' protocols, from task analysis, from interviewing subjects, from one's own intuition, from writing a simulation program, or from any other source. Each rule is written with the weakest preconditions possible so that it will be applicable in the widest possible range of situations. The second step is to fit this vocabulary of rules to the given protocol. At almost every point of the simulation, there will be many rules that can be applied because the rules have weak preconditions. The user selects one, and the simulation applies it. The user's job is to find a sequence of rule selections that maximizes the fit of the simulation to the protocol. If the protocol can be formally encoded, this step can be automated (VanLehn & Garlick, 1987; Kowalski & VanLehn, 1988). The result is a table which aligns the protocol with rule firings and "missed opportunities"—occasions where a rule could have fired but did not. For instance, in Tables 1 and 4 below, rows correspond to protocol lines and columns correspond to rules. Cells of the table contain a "1" if the rule fired during that line of the protocol, a "0" if it could have fired and did not, and a blank if it was not applicable at that time.

Given such a table, two kinds of analysis are performed. The first is simply to look for patterns in the firings

Table 1: Protocol and simulation of a subject solving a river crossing problem

Cycle	State	Protocol	Rule 1	Rule 2
1	LMSb, _	The boat can hold only 200 pounds? (E: The boat can hold only 200 pounds.) Okay... first... Small and Medium go back (E: Uh-huh.)... go across the river on it.	1	
2	L, MSb	And then, um, ... Oh... Large...	1	
3	"	[3 second pause](E: Yeah, go on... talk out loud.)...and... um... Large... um... [3 second pause] (E: Talk out loud. Tell me everything you're thinking.) But, I can't do it because someone has to sail the boat back.		0
4	"	(E: Ok... That's right. Somebody has to sail the boat back.) Oh! Ok... so... [4 second pause] Small sails the boat back and gets off,		1
5	LSb, M	and lets Large sail the boat back.	1	
6	S, LMSb	(E: Um-hmm. And then what happens.) Uh...	1	
7	"	[3 second pause] (E: Talk out loud.) And then Small... Small... can't think of anything... (E: Keep talking.)		0
8	"	So... Medium... sails back.		1
9	MSb, L	And... Medium and Small sail back.	1	
	LMSb	(E: Keep talking.) And they're all across! (E: Very good!)		

and missed opportunities. Using this kind of analysis, it was discovered that rules are rarely learned completely in one trial. Typically, the initial firing of a rule is followed by one or more missed opportunities, then another firing, a missed opportunity, and so on with an increasing ratio of firings to missed opportunities. Such patterns have theoretical implications. For instance, gradual acquisition is not consistent with the operation of learning methods, such as explanation-based learning (e.g., DeJong & Mooney, 1986), that acquire rules in one trial. In one case, the precise pattern of gradually increasing frequency of usage was predicted by a learning mechanism based on syntactic generalization of the rule.

Although the mere pattern of rule usage has shed some light on learning mechanisms, a second type of analysis has proved to be even more productive. This type of analysis examines the subject's utterances in the vicinity of each learning event. For instance, several authors have claimed that all learning occurs at *impasses* (Laird, Rosenbloom, & Newell, 1986; VanLehn, 1988). An impasse can only be precisely specified relative to a given simulation model, but the rough idea is that the model "doesn't know what to do next." If this idea is taken at face value, and impasse-driven learning is universal, then the subject's protocol at the first use of a rule should show signs of confusion or hesitation because the subject is at an impasse. The analyses presented below show that most initial uses of a rule are, as predicted, accompanied by unusually long pauses or by comments such as "I can't do it," or "It's not that easy." However, several cases were found where there are no such signs of an impasse at all. Thus, the data are consistent with the claim that most but not all rule acquisitions are triggered by impasses.

Study 1: A river crossing puzzle

The protocol analyzed in this initial study was not collected with the intention of studying the rule acquisition process, so it has some methodological flaws. However, it has the expositional advantage of being a very short protocol that nonetheless demonstrates some of the paper's claims.

The subject, a 9-year old girl, was given standard instructions for talk-aloud protocols then asked to solve the following puzzle:

Three men want to cross a river. They find a boat, but it is a very small boat. It will only hold 200 pounds. The men are named Large, Medium and Small. Large weighs 200 pounds, Medium weighs 120 pounds, and Small weighs 80 pounds. How can they all get across? They might have to make several trips in the boat.

The subject's protocol and an analysis of it appear Table 1. The first column numbers the cycles of rule execution. The second column abbreviates the puzzle's state--the notation "LMSb,_" means that Large, Medium, Small and the boat are on the source bank, and nothing is on the destination bank. The third column contains the subject's protocol. The remaining columns indicate rule firings and missed opportunities. The major rules used in the simulation appear in Table 2. Rule 1 is selected for firing at every opportunity (cycles 1, 2, 5, 6 and 9). The firing of the rule at cycles 2 and 6 causes a subgoal to be generated because the boat is not on the source bank. This assumes a cognitive architecture that automatically translates precondition violations into subgoals. (The simulation was conducted on Teton (VanLehn & Ball, in press), but ACT*, Soar, GPS and many other architectures have this

Table 2: Majors rules used in simulating the solution of the river crossing puzzle

1. If the goal is to move some men from bank X to bank Y, and the boat is on bank X, then load the boat to its maximum capacity, sail it across, and unload it.
 2. If the goal is that the boat be on bank X, and the boat is on another bank, Y, and the boat requires someone to sail it, and there are some men on bank Y, then load the boat with a small man, sail it across, and unload it.
-

property.) The subgoal of getting the boat back to the source bank causes the application of rule 2 at cycles 4 and 8.

Although all applications of rule 1 go smoothly, the first application of rule 2 is preceded by verbal evidence of an impasse. In the last line of cycle 3, the subject says, "But, I can't do it because somebody has to sail the boat back." The subject says she is stuck while at the same time mentioning an operation that could be applied. My interpretation of this line goes as follows. The subject recognizes that the puzzle situation is an idealization of reality, but she is unsure about how much of an idealization it is. In particular, she does not know whether the experimenter intends her to adopt the real-life constraint that most boats require a helmsman. The puzzle instructions do not state this constraint. (Indeed, another subject chose the other interpretation of the instructions, allowed the boat to sail itself back, and answered that it takes only two trips to get all the men across.) On this interpretation, the subject already "has" rule 2, but she does not know whether she is supposed to use it in this puzzle. After the experimenter confirms that this type of boat requires someone to sail it, the subject applies the rule (cycle 4). Although this learning event is arguably not rule acquisition, it is clearly learning of some kind. As will be seen shortly, it shares several properties with cases that are quite clearly rule acquisition events.

The second application of rule 2 is also preceded by signs of an impasse (cycle 7). The subject again claims to be stuck, saying "...can't think of anything...." Apparently, whatever she learned during cycle 3 is not immediately applicable at cycle 7. As will be seen later, this is quite typical--the first firing of a rule is followed by one or more missed opportunities. Several learning mechanisms are consistent with this behavior. For instance, it could be that the rule (or assertion) learned at cycle 3 has such a highly specific precondition that it does not apply at cycle 7, so the learning mechanism must create a generalized version of it. Another possible explanation rests on context effects--when the new rule is stored in memory, it is indexed in part by the context of cycle 3, which is assumed to be so different from the subsequent retrieval context (cycle 7) that retrieval fails. Although protocol data can differentiate such learning mechanisms (see below for an example), this particular case is consistent with a variety of learning mechanisms.

Depending on which learning mechanism one believes in, cycle 8 represents either a second firing of rule 2 or the firing of a new rule that is a generalization of rule 2. As mentioned earlier, this research method uses rules with the weakest preconditions possible, so in the simulation of this protocol, cycle 8 is a second firing of rule 2. However, this is just a methodological device for locating learning events in protocols. It is not intended as a claim about detailed learning mechanisms. I will continue to speak as if general rules were firing intermittently, even though it may be each general rule is actually an evolving collection of specific rules.

In summary, this protocol shows three interesting features. (1) The initial acquisition of the rule did not suffice to make it reliably operational. A second learning event was required (cycle 7). The second learning event took less time than the first (cycle 3). This pattern -- an initial formulation of a rule followed by one or more refinements of it -- occurs in later protocols as well. (2) Both learning events seem to be triggered by an impasse--a point where the subject does not know what to do next. Impasse-driven learning has been touted as a universal method for acquiring rules (VanLehn, 1988; Laird, Rosenbloom, & Newell, 1986). It will be seen later that although it is common for impasses to trigger learning, other types of triggering can also occur. (3) The subject reported neither the rule that was formed nor the processes that constructed it. The existence of the rule can be inferred from her actions, but if it was ever present in her working memory, she chose not to mention it. This too will turn out to be a common feature of learning events.

Study 2: The Tower of Hanoi

This study is a reanalysis of the classic protocol of Anzai and Simon (1979) wherein the subject invents several solution strategies for the five-disk Tower of Hanoi over the course of 90 minutes. During this time she receives no instruction. This corrects a methodological flaw in the first study, where the experimenter's comments seem to have been instrumental in the subject's learning.

Anzai and Simon uncovered the major strategies that the subject acquired and postulated learning mechanisms sufficient to acquire those strategies. They did not attempt a line-by-line comparison of the protocol and the behavior of their model. Using Teton (VanLehn & Ball, in press), we found that additional assumptions were necessary in order to achieve a line-by-line simulation of the protocol. The most important new assumption is that the subject has a strategy that develops around the first ten minutes of the protocol and remains fixed throughout the rest of the protocol. However, this strategy gives ambiguous advice on 25% of the moves, which will be called the *major* moves. On the major moves, the subject's fixed strategy narrows the choices down to two possibilities, but it does not say which one to take. The bulk of the subject's learning consists of a progression of strategies for making these decisions. With this new assumption, a model was formulated that fits the lines of the protocol with high accuracy, exceeding even that found in Newell and Simon (1972). The details can be found in VanLehn (1989).

Table 3 shows the rules used to make the major move decisions. Table 4 shows the analysis of the protocol. Each row of Table 4 is a major move. The first column numbers the major moves, the second column abbreviates the puzzle's state just prior to the move, and the third column abbreviates what the subject said while making the move. The puzzle's pegs are labeled A, B and C, where A is the peg that the disks start on and C is the peg they should end up on. The disks are labeled 1 through 5, with 5 being the largest disk. The notation "125,34,_" means that disks 1, 2 and 5 are on peg A, disks 3 and 4 are on peg B, and peg C is empty. The notation "2B, 1A" means that the subject announced a goal of moving disk 2 to peg B, then announced a movement of disk 1 to peg A. The notation "4pC" indicates a goal of moving a pyramid or group of four disks. Sometimes the subject announces a series of goals, pauses, and announces a different series of goals. This behavior is indicated by placing two rows in the table, one for each series of goals, and placing ditto marks in the first two cells of the second row. Horizontal lines in the table indicate places where the subject reset the puzzle to an initial state. The rightmost six columns of Table 4 show the applicability of rules. As always, a "1" indicates a rule firing, a "0" indicates a missed opportunity, and a blank indicates that the rule was not applicable. A "?" indicates that the rule may or may not have been fired -- the protocol evidence is unclear. The asterisks will be explained in a moment.

Table 3: Abbreviations and descriptions of rules for handling major moves

Initial rules

- *Look* The Anzai and Simon look-ahead search strategy.
- *1 blk* If the goal is to move a disk from one peg to another, and there is a single disk blocking the move, then get the blocking disk to the peg that is not involved in the move.
- *2 blk* If the goal is to move a disk from one peg to another, and the 2-high pyramid (i.e., disks 1 and 2) blocks the move, then get disk 1 to one of the pegs involved in the move (thus freeing disk 2 to move to the peg not involved in the move).

Rules acquired during the protocol

- *4B* Before attempting any of the top level goals, try to get disk 4 to peg B.
 - *Dsk* (The Anzai and Simon disk subgoal strategy.) If the goal is to get a disk from one peg to another, and there are some disks blocking the move, then get the largest blocking disk to the peg that is not involved in the move.
 - *Pyr* (The Anzai and Simon pyramid subgoal strategy.) If the goal is to move a pyramid from a peg to another peg, then get the next smallest pyramid to the peg that is not involved in the move.
-

Kurt VanLehn

Table 4: Rule firings and missed opportunities for the major moves

Move	State	Protocol	Initial rules			Learned rules		
			Look	1blk	2blk	4B	Disk	Pyr.
1	12345, ., .	1B	1			0	0	0
2	45, 3, 12	2B, 1A	1	1	?	0	0	0
3	5, 123, 4	5C, 1A	1			*1*	0	0
4	12345, ., .	1C	1			0	0	0
5	45, 12, 3	4B, 1A	0		1	*1*	0	0
6	5, 4, 123	1B	1			0	0	0
7	125, 34, .	1C	0		1	0	0	0
8	., 1234, 5	1A	1			0	0	0
9	3, 4, 125	1B	0		1	0	0	0
10	123, ., 45	3C, 1C	0		1	0	0	0
11	12, ., 345	1B	0	1		0	0	0
12	1, ., .	1C						
13	12, ., .	1B	0	1		?		
14	123, ., .	3C, 2B, 1C	0		0	*1*	0	0
15	., 12, 3	quits	0	1		0	0	0
16	1234, ., .	3B, 2C, 1B	0		0	1	0	0
17	4, 3, 12	1A	0		1	0	0	0
18	., 123, 4	1C	0		1	0	0	0
19	., ., 34	3B, 2A, 1C	0		0	*1*	0	0
20	12, ., 34	1B	0	1		0	0	0
21	12345, ., .	5C, 4B, 3C, 2B, 1C	0		0	0	1	0
22	45, 12, 3	4B, 1A	0		1	1	0	0
23	., ., .	4B, 2C, 1A	0		0	1	*1*	0
24	5, 4, 123	5C, 3B, 2A, 1A	0		1	1	0	0
25	125, 34, .	1C	0		1	0	0	0
26	., 1234, 5	4C, 3A, 2B, 1A	0		0	1	0	0
27	3, 4, 125	1B	0		1	0	0	0
28	123, ., 45	3C, 2B, 1C	0		0	1	0	0
29	12, ., 345	1B	0	1		0	0	0
30	12345, ., .	5C, 4B, 3C, 2B, 1C	0		0	0	1	0
31	45, 12, 3	4B, 1A	0		1	1	0	0
32	5, 4, 123	3pB, 1B	0		0	0	0	1
33	125, 34, .	1C	0		1	0	0	0
34	., 1234, 5	4pC, 4C, 3pA, 3A, 2C, 1A	0		0	1	1	1
35	3, 4, 125	2pA, 1B	0		0	0	0	1
36	123, ., 45	3pC, 2pB, 1C	0		0	0	0	1
37	12, ., 345	2pC, 1B	0	0		0	0	1

Some interesting findings are visible in the patterns of firings and missed opportunities. As in the river crossing puzzle, it is never the case that a rule is used consistently after it is first acquired. Instead, the usage of a new rule increases gradually. A second observation is that this subject occasionally compares the results of an old rule with those of the rule that supplants it. This can be seen in both cases where the subject re-does the planning of a move (moves 18 and 21). It can also be seen in move 32 where the subject mixes pyramid goals with disk goals. Thus, we do not see a rapid transition from an old rule to a new one, but a gradual transition that is sometimes accompanied by deliberate comparison of the old and new rules.

Obtaining further insight into the character of the rule acquisition process requires examining the protocol in the vicinity of the initial occurrences of the rules (see VanLehn, 1989, or Anzai & Simon, 1979, for the protocol itself). In the river crossing study, there were signs of impasses at both the initial firing of rule 2 and the subsequent firing. In this study, impasses were also present at most of the early rule firings. Asterisks are used in Table 4 to mark rule firings that were accompanied by long pauses and negative comments, such as "It's not that easy" or "I should have moved 5 to C." Impasses were common in the acquisition of both rule 4B and the disk subgoal rule. However, there seem to have been no impasses involved in the learning of the pyramid subgoal rule. At the first firing of the rule (move 30), the subject simply started phrasing her goals in terms of pyramid instead of disks. Instead of saying "3 will have to go to B..." as she said at move 22, she said, "I only need move three blocking disks to...B." There seems to have been no impasse here. At 32, the subject said

I will move the remaining four from B to C... It's just like moving four, isn't it? So... I will have to move 4 from B to C... For that, the three that are on top have to go from B to A... Oh, yeah, 3 goes from B to A! For that, 2 has to go from B to C, for that, 1 has to go from B to A.

Kurt VanLehn

Although this segment is long, it contains none of the signs of consternation that mark the other learning events. Instead, the subject seems to have been excitedly comparing the disk rule and the pyramid rule and proving to herself that they generate the same plan. If an impasse is defined to be an occasion when the subject does not know what to do, then this segment is not an impasse, for the subject seems to have two alternatives and believe that both are equally correct. In short, it seems that most rule acquisitions (2 of 3) are triggered by impasses, but rules can sometimes be learned without impasses.

There is a subtle pattern in the acquisition of the disk subgoal rule. Some of the early firings of the rule are marked by pauses and other signs of impasses, and some are not. Although space does not permit a detailed examination of the data (see VanLehn, 1989), it appears to be the case that the subject's initial formulation of the disk rule is highly specific in that it mentions the particular disks and pegs involved in the major move where it is acquired. Subsequent applications of the rule cause the names of specific pegs and disks to be replaced by variables. This gradual generalization of the rule means that some major moves can be handled by the evolving rule, while others cannot and force the rule to be further generalized. Pauses and other signs of impasses correlate perfectly with the places where generalization is predicted to occur. In particular, if it is assumed that the subject follows the policy of generalizing just enough to get the rule to accommodate the present situation, then it will take four learning events to learn a fully general version of the disk subgoal rule. All four of these predicted learning events are marked by impasses in the protocol (two occur during move 14, and both are marked by a distinct pause). So it appears that impasse-driven syntactic generalization, which has played an important role in several models of skill acquisition (e.g., VanLehn, 1983, 1986, in press), seems to be behind the acquisition of the disk subgoal strategy.

A last point to mention is that rules were discovered at a rate of about one every half hour (three were discovered in the 90 minute protocol). This rate seems to hold in the next study as well.

Study 3: College Physics

The protocols for the third study come from a study by Chi, Bassok, Lewis, Reimann and Glaser (in press) of eight students learning college physics from a standard textbook. Chi et al's study used a training format that comes close to the way students learn physics in college, except that the subjects could only refer to a textbook; they could not ask questions of a teacher. The subjects first learned the initial four chapters of a standard college physics textbook to criterion. They then read the fifth chapter, which covers the target subject matter, Newtonian particle dynamics. When they came to the worked examples at the end of the chapter, protocol collection began. Protocols were collected as the students studied 3 examples and worked 19 problems. The examples and the problems present ample opportunities for learning because they address issues that simply are not covered anywhere in the preceding material. For instance, the concept of a "normal force" is first introduced in the context of an example. The students took between 8 and 29 hours to complete the study. The protocols cover the last 3 to 6 hours.

As simulations are currently being constructed for each of the 8 protocols, it is too early to report accurate data on learning events. However, Bernadette Kowalski and I could not resist doing a hand analysis of one protocol. We found clear indications of five rules being acquired. As the protocol lasts 3.5 hours, this is an average of one rule every 40 minutes, which is comparable to the rate found in the Tower of Hanoi study (one rule per 30 minutes). We found some evidence that rules are acquired gradually, but we are reluctant to put a number on it because it is difficult to detect missed opportunities without a simulation. The usual signs of impasses marked the initial firing of 3 of the 5 rules. The other two rules seem to be acquired as the subject reflects on a just-completed solution. As an illustration, the next few paragraphs present one of these rule's acquisition event.

Subject 101 is confused about the difference between weight and mass throughout most of the experiment. (Many other students had the same confusion.) Eventually he discovers that weight is the force due to gravity while he is solving the following problem: "A fireman weighing 160 pounds slides down a vertical pole with an average acceleration of 10 feet per second. What is the average vertical force he exerts on the pole?" The subject reads the problem, then says:

6. Okay. Um, we'd have to consider...the force of gravity.
7. Okay. Let's find out what the force of gravity is exerting
8. on them and then we can figure out what, what his, what he's
9. exerting on it.
10. Now, let me remember, weight is equal to, what's force equal
11. to, weight?
12. Force is equal to weight over gravity times acceleration.

Kurt VanLehn

The subject decides to follow a generic plan, which he has used many times before. The plan is to find the forces acting on the body (the fireman, in this case), sum them, and apply $F=ma$. He summarizes his intentions in lines 7, 8 and 9. However, the plan's first goal, which is to find the force of gravity acting on the fireman, thwarts him. He does not know that the 160 pound weight is the force of gravity on the fireman, so he sets about to calculate the force using the derived law, $F=(W/g)a$. After fumbling with the units and looking up the appropriate value for the gravitational acceleration, g , he substitutes the freefall acceleration for a and obtains $F=(W/g)g=W$. At this point, he says:

- 42. Oh, I'm going to get force is equal to weight divided by
- 43. gravity times gravity which is going to be equal to weight.
- 44. Right?
- 45. Is that right?
- 46. Okay. Um, so I'm going to get 160 pounds. That's the force.
- 47. Yeah. It kind of makes sense 'cause they, they weight you in
- 48. pounds, don't they?
- 49. That's force.
- 50. Okay. So, average acceleration, the force he, the gravity is
- 51. exerting on him, yeah, yeah, that makes sense, is 160.

At line 43, he has the solution to his subgoal. He double checks the math in lines 44 and 45 (probably), and again states the solution in line 46. Although he could simply go on to the next step in his plan, the simplicity of the equation $F=W$ apparently prompts him to reflect on his solution. Thus, a learning event begins around line 47. The subject appears to use a kind of explanation-based reasoning. Although he has just built a proof that $F=W$ for this problem, he adds a second "proof" based on the units of force and weight (both are measured in pounds). This seems to be critical to establishing the generality of the result, which is that weight is the force exerted by gravity on an object. The learning event ends at line 49, and the subject returns to the plan in line 50. However, he indulges in one last check of the result, in line 51, before going on to finish the problem off. The next time he has an opportunity to apply his new rule, he initially fails to retrieve it, but is reminded of it halfway through the problem, and happily applies it. Thereafter, he always uses the new rule whenever it is applicable.

This segment of the protocol illustrates that rule acquisition in a knowledge-rich context has much the same character as it does in the knowledge-lean context of learning to solve a puzzle. For instance, it appears that the acquired rule is not completely learned during the first instance of its use, for the subject nearly misses the opportunity of applying it later. This particular rule does not seem to be acquired at an impasse. Instead, the subject seems to infer it while reflecting on his just-completed solution to a subproblem (line 47). However, other physics rules do seem to be learned at impasses.

Conclusions

Three analyses have been presented showing that the initial uses of problem solving rules can be located in protocol data. This analysis method yielded the following observations about the acquisition of rules:

1. Rules are seldom completely learned in one trial. The initial firing of a rule is often followed by several missed opportunities before the rule comes to be fired at every opportunity.
2. Sometimes, this gradual increase in applicability is consistent with a learning mechanism (VanLehn, 1983, in press) that operates by initially constructing a highly specific rule then generalizing it only when an impasse forces it to.
3. Long pauses, negative comments and other signs of impasses are common at the early firings of rules, but some rules are acquired without any visible signs of an impasse.
4. Sometimes, the subject explicitly compares a new rule to the old rule that it replaces. This indicates that the subject is probably aware of both of them, although none of the subjects in any of these studies explicitly mentions or describes their rules. It also indicates a more-or-less deliberate application of a method for improving one's knowledge.
5. In the context of learning-by-doing, wherein the subject receives no instruction from tutors or manuals, rule acquisition occurs at the rate of about one rule every half hour.

Acknowledgments

I would like to thank Bill Ball and Bernadette Kowalski for their indispensable help with the analysis, and Micki Chi for her thoughtful advice. This research was supported by the Cognitive Sciences Division and the Information Sciences Division of the Office of Naval Research under contracts N00014-86-K-0678 and N00014-88-K-0086.

References

- Anderson, J. R., Farrell, R., & Saurers, R. (1984). Learning to program in LISP. *Cognitive Science*, 8, 87-129.
- Anzai, Y. & Simon, H.A. (1979). The theory of learning by doing. *Psychological Review*, 86, 124-140.
- Chi, M.T.H., Bassok, M., Lewis, M., Reimann, P. & Glaser, R. (in press, 19??). Learning problem solving skills from studying examples. *Cognitive Science*, .
- DeJong, G. & Mooney, R. (1986). Explanation-based learning: An alternative view. *Machine Learning*, 1(2), 145-176.
- Klahr, D., Langley, P. & Neches, R. (1987). *Production System Models of Learning and Development*. Cambridge, MA: MIT Press.
- Kowalski, B. & VanLehn, K. (1988). Inducing subject models from protocol data. In V. Patel (Eds.), *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Laird, J. E., Rosenbloom, P. S., and Newell, A. (1986). Chunking in Soar: The anatomy of a general learning mechanism. *Machine Learning*, 1(1), 11-46.
- Newell, A. & Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.
- VanLehn, K. (1983). Human skill acquisition: Theory, model and psychological validation. In *Proceedings of AAAI-83*. Los Altos, CA: Morgan Kaufmann.
- VanLehn, K. (1988). Toward a theory of impasse-driven learning. In H. Mandl & A. Lesgold (Ed.), *Learning Issues for Intelligent Tutoring Systems*. New York, NY: Springer Verlag.
- VanLehn, K. (1989). *Learning events in the discovery of problem solving strategies* (Tech. Rep. PCG-17). Dept. of Psychology, Carnegie-Mellon University.
- VanLehn, K. (in press, 19??). *Mind Bugs: The origins of procedural misconceptions*. Cambridge, MA: MIT Press.
- VanLehn, K. & Ball, W. (in press, 19??). Teton: A large-grained architecture for studying learning. In VanLehn, K. (Ed.), *Architectures for Intelligence*. Hillsdale, NJ: Erlbaum.
- VanLehn, K. & Garlick, S. (1987). Cirrus: an automated protocol analysis tool. In Langley, P. (Ed.), *Proceedings of the Fourth Machine Learning Workshop*. Los Altos, CA: Morgan-Kaufmann.